

## Minimum-resource distributed navigation and mapping

Douglas W. Gage

Space and Naval Warfare Systems Center San Diego  
SPAWARSYSCEN D371, San Diego, CA 92152-7383

### ABSTRACT

This paper proposes a conceptual design for a distributed system of very simple robots capable of performing a useful real-world mission such as mapping the interior of a building overnight with a swarm of possibly hundreds of cockroach-sized robots. Presentation of this system concept follows an initial discussion of strategies for developing distributed robotic systems. Success is dependent on making good decisions in selecting appropriate applications, in system design, and in executing the system development process.

Each robot includes basic mobility, crude odometry, contact or near-contact object/obstacle detection sensors, an omnidirectional beacon (probably IR), and a beacon detection sensor that can simultaneously detect multiple beacons on other robots and measure the bearing of each to less than one degree. Beacon triangulation (combined with knowledge of some baseline distance) allows the determination of the position of any robot (and any object next to it) relative to the others. Occlusion of a robot's beacon indicates the presence of an intervening object, while lack of occlusion identifies a "ray" of free space. Clever deployment of large numbers of robots will permit mapping of walls and objects, and can also support the detection and localization of intruders moving within the space. The object sensor can be very short range and, therefore, hopefully very simple and cheap, perhaps implemented as an array of whiskers. The beacon sensor is more complex, but it can be completely tuned to the detection of the cooperating beacon. Thus, the robots need tackle no perception tasks whatsoever.

The system development begins with the implementation of an initial baseline system, in which each robot is controlled by a central coordinating element via a high bandwidth communications link. This initial effort will develop and validate behaviors and algorithms, assess sensitivity to sensor error, determine communications and processing requirements, and generally expedite the system's evolution into a fully distributed system.

**Keywords:** distributed, robots, navigation, mapping, beacon

### 1. INTRODUCTION

The notion that there are some tasks that can be performed most effectively by a large number of small inexpensive mobile robots<sup>1</sup> has developed considerable currency in the robotics community. Motivated in part by the perceived "industry" of ant or bee colonies, a body of research<sup>2</sup> exploring the issues associated with creating and exploiting such "swarm" systems has been developed over the past 10 to 15 years. In addition, "nano" scale robots have found a constituency in the popular culture, as exemplified by the novel *The Diamond Age*<sup>3</sup>. To date, however, no distributed robotic systems have been deployed, and no physical systems of more than 20 to 30 mobile robots have been demonstrated, even in a research environment.

The goal of this paper is to sketch out the design of a distributed robotic system that can first be developed and then be mass-produced at the lowest possible cost and with the lowest possible risk, but which will still be capable of reliably performing useful tasks in the real world. The underlying assumptions of the exercise are that (1) the process of implementing physical robots always turns out to be more difficult, to take more time, and to cost more than anticipated, and (2) a robot's ability to perceive important task-relevant features of its environment is unreliable at best and non-existent at worst, all the more so when its sensors must be compact, low-power, and inexpensive. This suggests that it would be wise to start by choosing an application domain that can be addressed by using only the simplest robotic hardware and the least challenging perception capabilities. Similarly, the system design and development process should lead to the minimum-hardware and minimum-perception solution that satisfies the application requirements.

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE <b>NOV 2000</b>		2. REPORT TYPE <b>N/A</b>		3. DATES COVERED <b>-</b>	
4. TITLE AND SUBTITLE <b>Minimum-Resource Distributed Navigation and Mapping</b>				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) <b>Douglas W. /Gage</b>				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) <b>Space and Naval Warfare Systems Center SPAWARSYSCEN D371 San Diego, CA 92152-7383</b>				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT <b>Approved for public release, distribution unlimited</b>					
13. SUPPLEMENTARY NOTES <b>The original document contains color images.</b>					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT <b>UU</b>	18. NUMBER OF PAGES <b>8</b>	19a. NAME OF RESPONSIBLE PERSON
a. REPORT <b>unclassified</b>	b. ABSTRACT <b>unclassified</b>	c. THIS PAGE <b>unclassified</b>			

Section 2 presents some of the important factors that should be considered in choosing (1) an application, (2) a system design, and (3) a development process in order to follow a "path of least resistance" toward the development of a practical distributed robotic system. Section 3 provides a brief discussion of the specific application domain of tactical military operations inside buildings. Then Section 4 presents a specific distributed robotic system concept for indoor navigation and mapping intended to approximate a "path of least resistance" approach to address the interior reconnaissance and surveillance application.

## **2. STRATEGIES FOR APPLICATION SELECTION, SYSTEM DESIGN, AND THE DEVELOPMENT PROCESS**

### **2.1 Strategies for Application Selection**

Research to date in distributed robotics has been driven principally by technology push. While some of the basic research and exploratory development work that has been pursued has addressed specific application areas, demonstrations of capabilities seen to date have not justified the initiation of a full scale program to develop a system suitable for production and deployment. In order to reach this point, it is necessary to identify an important and useful real world application for whose requirements a distributed robotic system offers not just a feasible solution, not just a good solution, but indeed the best solution or perhaps even the only solution. Here are some of the characteristics that such an application might possess:

**Opportunity to exploit parallelism for speed.** Performing multiple operations in parallel is a classic strategy for speeding up a process. However, this will provide a convincing rationale for using distributed robotics only if (a) speed is a key requirement, (b) the process is fully parallelizable, and (c) other possible competing strategies for achieving the required speed are not as effective or efficient.

**Requirement for intrinsic parallelism.** Some processes require simultaneous action at different places. For example, while a search for stationary targets can almost always be accelerated by using multiple searchers, a search for targets capable of movement may require the use of multiple searchers in order to prevent the targets from evading detection indefinitely.

**Requirement for expendability.** If an application requires the repetitive performance of a task that could result in the performer's being disabled or destroyed, then using multiple devices is necessary to ensure that the mission can be completed successfully. An example is the PUCA (Pick Up and Carry Away) strategy for cleaning up unexploded cluster munitions (UXO), as addressed in NAVEODTECHDIV's Basic UXO Gathering System (BUGS) Program<sup>4</sup>, part of OSD's Joint Robotics Program.

**Requirement for "dedication".** Some applications may require that a device be physically allocated (dedicated) to each of a (potentially large) number of specific places determined by environmental factors. One example is the BIP (Blow In Place) strategy for cleaning up UXO (the simultaneous detonation of explosive charges placed next to each UXO). Other examples include building a network of communications relays, providing continuous detection of intruders at multiple physically separated possible entry points, or simply picking up the four corners of a large object which is to be moved.

**Minimal ancillary technical challenges.** Some applications that would otherwise be good applications for distributed robotic system solutions require solving very difficult technical challenges specific to the application. For example, the clearance of anti-personnel (AP) landmines is clearly a problem of widely acknowledged importance, and it exhibits both the expendability and dedication requirements described immediately above. The problem is that we have no sensor that can reliably detect the presence of a buried AP landmine, much less a sensor small enough and inexpensive enough to be incorporated into a distributed robotic system.

**Coherent user view of distributed system functionality.** A system must provide its users with a well defined and comprehensible view of system-level functionality: how to task, monitor (and override, if necessary), and assess the results of the system's operation – all without having to be aware of the specific activities of individual robots. For some highly distributed system applications, this may be extremely difficult to realize. For many applications, such as search, the coverage paradigm<sup>5</sup> may provide a useful metaphor.

These technical characteristics, however, can only serve as basic guides for the selection of a target application. Especially for the military, the decision to fund the expensive development of a robotic system is an administrative process that involves political and economic factors as well as technical ones. But the probabilities of success are maximized with an application that makes sense technically and user proponents who understand the proposed system's capabilities and the roles it can play in enhancing mission effectiveness.

## 2.2 Strategies for System Design

Just as it is possible to identify technical indicators that an application might be a strong candidate for a distributed robotic solution, it is also possible to identify a number of technical characteristics that would be highly desirable in a distributed robotics system design:

**Scalability in numbers.** A systems approach that is capable of supporting extremely large numbers of robots is desirable, to avoid arbitrary limits on the size of applications that can be handled. Command, Control and Communications (C3) architectural considerations will usually be the critical determinants in this dimension.

**Scalability in size.** Making robots smaller can increase their operational effectiveness by, for example, allowing entry into tight spaces and supporting covert operations through reduced detection signatures. The system design should therefore not require either physically large antennas or active sensors that consume a lot of power. In addition, making the robots smaller makes them easier to deploy in quantity and may possibly permit reduced manufacturing costs.

**Low-cost manufacturability.** It is clearly desirable that the robots be designed so that manufacturing economies of scale can make the system affordable in production.

**Minimized ancillary technical challenges.** It is important to focus scarce development resources on the work that absolutely has to be done. The system design should avoid the use of immature subsystem technologies or components that involve technical, schedule, and cost risks that could otherwise be avoided.

## 2.3 Strategies for the Development Process

The track record for development programs aimed at producing autonomous mobile robot systems or unmanned ground vehicles (UGVs), going back to DARPA's Autonomous Land Vehicle Program more than fifteen years ago, is one of less than overwhelming success. While there are a variety of reasons for this, only some of them technical, it is possible to identify some important lessons learned from these experiences:

**Focus initial attention on the robots' interactions with the environment.** Most robotic vehicle programs have been challenged (if not derailed) by the inadequacies of available perceptual capabilities<sup>6</sup>. Even if a sensor provides accurate readings, these readings may not provide the answers to the questions that need to be answered. For example, it's probably more important for a vehicle's navigation process to know whether the object detected in its path is a boulder instead of a tumbleweed than whether its distance ahead is 10.2 instead of 10.4 meters.

**Leverage the resources of other programs.** Especially in urban terrain, communications is a major challenge. Fortunately, other programs are spending a lot of money learning to network manned vehicles, soldiers, command posts, etc. together. Robotic systems, when ultimately deployed, should use these same communications solutions whenever possible. Similarly, enhanced power technologies, fuel cells as well as improved batteries, are being aggressively pursued by others. A robotics program needs to focus its resources on solving the problems unique to robotic systems.

**Decompose the system.** It is critical to identify the various elements of the system, to establish the functional and performance requirements for each subsystem, and to precisely specify the interfaces between them. Understanding requirements at the subsystem level enables effective leveraging of other programs' work and also serves to identify specific critical missing capabilities.

### 3. TACTICAL MILITARY OPERATIONS INSIDE BUILDINGS

As the canonical strategic military threats of the Cold War era have receded, greater interest has been focused on other military operations, many of which include activities in and around buildings (referred to as MOUT, or Military Operations in Urban Terrain). DARPA ATO's Tactical Mobile Robotics (TMR) Program<sup>7</sup> has addressed the development of mobility, navigation, and mapping<sup>8</sup> technologies to support the operation of backpack-sized and smaller "throwable" robots inside buildings. This section discusses the characteristics of two classes of military operations that take place inside buildings as potential applications for distributed robotic systems.

One of the archetypal MOUT operations is that of "clearing" a building, in which a team of dismounted warfighters moves systematically through each of the rooms on each of the floors of the building, eliminating any threats that they encounter. This is dangerous work, and there would be an obvious payoff to using robots to reduce the human warfighter's exposure to danger during this process. However, it is clear that current robots can not perform this operation without human participation, and warfighters have raised a number of potential objections to working alongside robots. For example:

- When warfighters have the element of surprise, the robots must not compromise it.
- Warfighters must never be forced to slow the tempo of their operation by having to wait for the robots.
- Since a warfighter's physical survival depends on staying constantly alert to possible threats, robots must never provide a deadly distraction.

These requirements pose extremely difficult challenges to the robotic system developer. However, a second possible application of robots inside a building is to perform covert reconnaissance independent of any sort of "clearing" operation. This function has relevance to law enforcement as well as military special forces operations. Compared to building clearance, the reconnaissance application favors a distributed robotic solution in several ways:

- Reduced speed requirement for the robot platforms, since they are not working in concert with personnel.
- Possibly reduced requirement for traversing rough terrain because of the absence of battle damage – but ingress and egress may still present critical mobility challenges such as stairs.
- Opportunity for more detailed mission planning: a law enforcement or special forces scenario may afford more preparation time, better trained operators, and more complete information about the area of operations.
- Increased importance of covertness, therefore increasing the value of employing smaller platforms that tend to be naturally harder for the adversary to detect than larger robots or humans.

While building reconnaissance seems to be superior to building clearance as an application, it does not score particularly well in terms of the criteria presented in Section 2.1. Since it involves no requirements for intrinsic parallelism, expendability, or dedication, it is almost certainly not the "Killer App" needed to revolutionize the world of distributed robotics. On the other hand, building reconnaissance does exploit parallelism to speed up the mapping process, and, perhaps most critically, it avoids the ancillary technical challenges presented to small robots by outdoor mobility. Moreover, reconnaissance is an application with recognized military significance. This paper therefore addresses an abstracted version of the building reconnaissance application. Disregarding any specific performance requirements, the focus is on the design of a distributed system of small robots that can navigate within and map one floor of a building and also provide some basic information about the presence of people moving about within this space.

### 4. A "MINIMUM-RESOURCE" SYSTEM FOR INTERIOR NAVIGATION

A conceptual design is presented here for a distributed robotic system capable of mapping building interiors. The goal is that this system should employ the "minimum-resource" robots capable of performing this task – that no other system could do this job while employing robots that would be cheaper than the ones proposed here (in their final evolved form). The simplicity of the system is achieved by completely eliminating the need for the detection or perception at a distance of any object, except for the beacons mounted on other robots (and, by implication, objects which occlude the beacon transmission between two robots). The geometric relationship between the robots is determined by triangulation, while contact sensors detect the presence of objects adjacent to the robots.

No claim is made that the function and form of this proposed system is unique or novel. There are obvious parallels to coastal piloting<sup>9</sup> (in the use of landmarks and navigational aids), aviation navigation (from beacon to beacon), traditional pre-GPS

surveying<sup>10</sup> (networks of triangulation stations), as well as robotic navigation using artificial landmarks. Specifically, navigation schemes that make use of triangulation among cooperating robots have been explored by Kurazume et al<sup>11,12,13</sup>. Moreover, CMU's modular "millibot" platform<sup>14</sup> could, with the addition of a module containing the IR beacon and beacon sensor, serve as a capable initial testbed vehicle to support the continuing evolutionary development of the system.

#### 4.1 The Initial Baseline System

We will first describe the initial baseline system, whose principal role is to support exploratory interaction of the robots with the environment, and then outline a system development process that will provide an evolution of this initial design to yield a fully distributed system offering enhanced performance and minimum cost, as described in Section 4.2.

The baseline system is comprised of tens to hundreds of identical ("single caste") robots. In the initial baseline implementation, the robots look like small toy cars, nominally about 3 to 4 inches long, and each robot incorporates the following capabilities:

**Mobility.** The ability to move over typical office floors, including wood, tile, and short carpet. Maximum speed capability is not important. The ability to steer left and right, but precision of steering is not important.

**Odometry.** The ability to measure distance traveled to within a few percent. Precision is not important. Measurement of actual vehicle movement, rather than of wheel or track motion, would be an advantage, since it would eliminate errors due to slippage.

**Object/obstacle detection sensors.** The ability to detect objects with which the robot comes into contact or near-contact, and to identify where the contact occurred (front, back, left, right). It would be a strong plus if the sensors provided 2 bits (2-4 levels) of distance resolution, to allow the vehicle to navigate along ("follow") a wall without having to repeatedly "bounce" off it.

**Omnidirectional beacon.** A beacon (probably IR) capable of being detected by the other robots in the area. Optics may be used to focus the beacon strength in the horizontal plane, but, while longer range is better, there is no critical range requirement. The ability to turn beacon transmission on and off.

**Beacon detection sensor.** The ability to detect multiple beacons of other robots in the horizontal plane (nominally plus or minus 15 degrees -- this value to be refined later) and support the measurement of their relative bearings with a precision of a fraction of a degree. This is basically a camera with a cylindrical field of view 1 pixel high by 512 or (better) 1024 pixels around. This sensor is the single critical "magic bullet" in the system.

**Beacon reflector.** Cylindrical mirror mounted coaxially with the beacon and beacon detection sensor, so that another robot will see its own beacon reflected at the bearing of this robot's beacon, if the two are close enough. This allows a robot to serve as a permanent short range beacon for navigational purposes, even after its batteries have died. In addition, the rough single-bit measurement of the distance to another robot provided by the limited detection range of the reflection should provide a useful input to various robotic behaviors.

**Radio transceiver.** The ability to provide bi-directional digital communications with the system's central controller. All robots transmit on the same channel, with each robot transmitting only when authorized by the central controller, which transmits either on the same channel as the robots or on another channel to which the robots' receivers are tuned. The emerging Bluetooth technology may provide candidate components.

**Onboard Processing.** Local processing sufficient to allow the central controller, via the radio channel, to control/measure vehicle speed and steering angle, read the object detection sensors, turn the beacon on and off, and read the beacon detection sensor. Ability to execute commanded behaviors such as wall following by using the object detection sensors.

**Power.** Appropriate off-the-shelf batteries.

Although virtually everything these very simple robots do will be in response to commands issued by the system's central controller, the system's physical interface to the environment will be the same as if the processing were actually implemented on the robots themselves. Functionally, therefore, the system should be capable of executing the distributed behaviors of interest, although performance limits due to communications and processing latencies will likely be severe and become more severe as the number of robots active at any time is increased. While this central controller could be implemented as a single monolithic process, essentially implementing a single organism with each of the robots serving as an appendage, it makes good sense to partition the controller into a set of "proxy" processes, one for each robot.

Building on low-level primitive commands, higher-level functionalities of interest will be explored, including:

**Localization of robots.** The primitive command for robot localization is "Triangulate", which takes as its argument a list of robot IDs. The robots involved must be stationary. Each robot in sequence turns on its beacon and acknowledges that it has done so (so that other robots can determine which beacon is which). Then each robot in sequence transmits to the central controller a list of the bearings of each of the other beacons it can see. Knowing some initial baseline distance, the central controller can now compute the geometry of the robots.

**Mapping of contacted objects.** The central controller can compute the approximate position of an object that a robot encounters with its contact sensors by knowing the position of the robot. The robot can be commanded to follow the perimeter of an extended contacted object (CO), and periodic measurements of its position will trace out this perimeter.

**Mapping of free space.** If another robot's beacon stays in continuous view to another robot, then the system knows that there are no Beacon-Occluding Objects (BOOs) along the line between them. As either or both of the robots move through the environment, they sweep out an area of known free space between them.

**Detection and mapping of beacon-occluding objects.** If a moving robot's view of a stationary robot's beacon becomes occluded, then it can back up slightly, reacquire the beacon, turn toward it, and then attempt to move toward it while intermittently moving into the beacon occlusion field of the object. In most cases, the robot's contact sensors will eventually encounter the object, and it can then proceed to map it as a contacted object. But, of course, as will be iterated below, not every BOO is a CO, or vice versa.

**Coordinated group movements.** While an individual robot in isolation may not be capable of traveling in a straight line, robust group behaviors (somewhat analogous to the "gliders" found in John Conway's game of Life) can be designed to implement a library of motion primitives that can support the systematic exploration of the environment.

**Dedication of robots.** Since, in this triangulation scheme, distance measurements are scaled to some initial baseline(s), maps created by using the techniques above will not be metrically precise. By maintaining some number of robots in key locations to serve as permanently positioned landmarks, it will be possible for other robots to effectively navigate using the map. Moreover, landmarks can be reliably reinstalled later in places with unique contact object configurations, such as room corners. Finally, the inclusion of a beacon reflector on each robot allows it to serve as a passive-but-detectable stationary landmark even after its battery has died.

Moreover, the system will be used to explore some of the inevitable complications of operation in the physical world, such as:

- Conjunction of beacons – the eclipse of one beacon by another.
- Reflection – both specular and diffuse, by objects of various sizes and shapes, and of a robot's own beacon or another robot's beacon – can generate indications of "ghosts" that must be interpreted correctly.
- The fact that an object that is detected by a contact sensor may not occlude a beacon – for example, it may not be tall enough. All COs are not necessarily BOOs.
- The fact that an object that occludes a beacon may not be detectable by a contact sensor. All BOOs are not necessarily COs.
- Dealing with the consequences of negative obstacles, including the sudden "disappearance" of robots that fall into holes.

## 4.2 The Evolution of the Target System

The purpose of the initial development phase is to guide the evolutionary development of a deployable target system. Specifically, exploratory experiments using the initial baseline system will serve to:

- Validate the functionality and assess the performance of the initial physical robot configuration in terms of mobility and sensing.
- Validate the functionality of the robots' software perceptual and behavioral algorithms in real world environments, and determine the dependency of algorithm performance on input sensor performance.
- Validate the interprocess coordination infrastructure and the protocols for inter-robot communications by implementing them in the context of the "proxy" processes representing each of the robots in the central controller.
- Determine communications bandwidth and latency requirements for the target robot system.
- Determine processing throughput requirements for the target robot system.
- Estimate power requirements for a fully distributed system.

The second phase of the development will be to implement a fully distributed system in which each robot will host its own onboard processing resources. The proxy software developed for the baseline will be ported to the robots themselves, and the central control unit will shrink to provide only system-level coordination, and to support operator tasking, monitoring, and override capabilities. The deferral of important subsystem-level design decisions to this second phase means that these decisions can now be fully informed by the results of the experiments using the initial baseline system. Processing, communications, and power resources can be sized to fit the actual needs of the application. In addition, subsystem components exploiting newer technologies may become available for inclusion in the system.

This second phase also provides the opportunity (and poses the necessity) of developing a much more capable and highly integrated version of the beacon system. Ultimately, an integrated module will incorporate the beacon transmitter (capable of software-controlled modulation), cylindrical reflector, and beacon detection sensor coaxially mounted in the same package, with optical elements that provide gain in the horizontal plane for both the transmitter and detector. Processing integrated with the annular photodetector FPA will allow the module to provide the robot's main processor with:

- The beam width and integrated intensity of each detected beacon, to provide some measure of its distance
- Each detected beacon's bearing and rate of change of bearing, to support triangulation

In addition, several modes of communication can be encoded in the beacon modulation, including transmitter ID, signboard/pheromone display, and broadcast messaging, as well as traditional connection-oriented communications services.

The beacon module requires no exotic component technologies. The principal task is to design and implement an integrated detector/processing chip, which should be easily realizable with standard CMOS fabrication technology. In large quantities, the manufacturing cost of the complete module should easily be less than \$10. Incorporation of the module into a mainstream toy system consisting of interacting mobile robots would easily justify the non-recurrent engineering costs required.

The fully distributed system implemented in this second phase will support the continuing evolution of the generic system infrastructure and the adaptation of the system to serve the specific needs of specific applications, which will in turn motivate future iterations of the system in which the robots are smaller and more tightly integrated.

## 5. CONCLUSION

The development of a distributed system of small inexpensive robots presents so many challenges that it makes good sense to adopt a "path of least resistance" approach. In summary, this means selecting a technically feasible application for which a distributed solution will provide a unique payoff, developing a system design that avoids unnecessary technical challenges, and executing a development process that focuses program resources on the critical problems that must be solved.

Absent the identification of a true "Killer App" for distributed robotics, the military problem of reconnaissance within building interiors has been addressed with a conceptual system design that aspires to perform indoor navigation and mapping in a "minimum resource" fashion. The key is to solve the robot localization problem by triangulation using beacons mounted on each of the robots, and to localize objects detected by contact sensors with respect to the adjacent robots. Areas of free space and areas containing beacon-occluding obstacles can be identified. No perception of objects at a distance is employed. A development strategy has been proposed in which an initial baseline system employing a central controller to control the



robots is used to identify and validate subsystem (e.g., communications, processing, power) requirements for the second phase fully distributed system. A highly integrated beacon module has been identified as the “magic bullet” needed to realize the implementation of this system.

## ACKNOWLEDGEMENTS

This work has been performed under the Software for Distributed Robotics (SDR) Program sponsored by the Defense Advanced Research Projects Agency (DARPA) Information Technology Office (ITO). The author acknowledges numerous productive conversations with many participants in the SDR Program, DARPA/ATO’s Tactical Mobile Robotics (TMR) Program, and DARPA/MTO’s Distributed Robotics (DR) Program,

## REFERENCES

1. A. M. Flynn, "Gnat Robots (And How They Will Change Robotics)", *Proceedings of the IEEE Micro Robots and Teleoperators Workshop*, Hyannis MA, 9-11 November 1987. Also appeared in *AI Expert*, December 1987, pp. 34 et seq.
2. L. E. Parker, "Current State of the Art in Distributed Autonomous Mobile Robotics", in *Distributed Autonomous Robotic Systems 4*, L. E. Parker, G. Bekey, and J. Barhen eds., Springer-Verlag Tokyo 2000, pp. 3-12.
3. N. Stephenson, *The Diamond Age*, Bantam Books, New York, 1995.
4. C. DeBolt et al, "Basic UXO Gathering System (BUGS): Multiple Small Inexpensive Robots for Autonomous UXO Clearance", *Proceedings of Third International Symposium on Technology and the Mine Problem*, Monterey CA, 6-9 April 1998.
5. D. W. Gage, "Command Control for Many-Robot Systems", *AUVS-92, the Nineteenth Annual AUVS Technical Symposium*, Huntsville AL, 22-24 June 1992. Reprinted in *Unmanned Systems Magazine*, Fall 1992, Volume 10, Number 4, pp. 28-34.
6. H. R. Everett, D. W. Gage, G. A. Gilbreath, R. T. Laird, and R. P. Smurlo, "Real-world Issues in Warehouse Navigation," *SPIE Mobile Robots IX*, Vol. 2352, Boston MA, November 1994, pp. 249-259.
7. J. G. Blitch, "The Tactical Mobile robotics Program," *SPIE Mobile Robots XIV*, Vol 3838, Boston MA, September 1999.
8. S. Thrun, W. Burgard, and D. Fox, "A Real-time Algorithm for Mobile Robot Mapping with Applications to Multi-Robot and 3D Mapping," *Proc. 2000 IEEE International Conference on Robotics and Automation*, San Francisco CA, April 2000.
9. C. F. Chapman, *Piloting, Seamanship, and Small Boat Handling*, 51<sup>st</sup> Edition, Hearst, New York, 1974.
10. R. E. Davis, F. S. Foote, J. EM. Anderson, and E. M. Mikhail, *Surveying Theory and Practice* (Sixth Edition), McGraw-Hill, New York, 1981.
11. R. Kurazume, S. Nagata, and S. Hirose. "Cooperative Positioning with Multiple Robots", *Proc. 1994 IEEE International Conference on Robotics and Automation*, Los Alamitos CA, 8-13 May 1994, volume 2, pp. 1250-1257.
12. R. Kurazume, S. Hirose, S. Nagata, and N. Sashida. "Study on Cooperative Positioning System (basic principle and measurement experiment), *Proc. 1996 IEEE International Conference on Robotics and Automation*, Minneapolis MN, 22-28 April 1996, volume 2, pp. 1421-1426.
13. R. Kurazume, and S. Hirose. "Study on Cooperative Positioning System: Optimum Moving strategies for cps-iii", *Proc. 1998 IEEE International Conference on Robotics and Automation*, Leuven Belgium, 16-20 May 1998, volume 4, pp. 2896-2903.
14. [http://www.ri.cmu.edu/projects/project\\_343.html](http://www.ri.cmu.edu/projects/project_343.html)